

EXHIBIT A

**FEDERAL TRADE COMMISSION (“FTC”)
CIVIL INVESTIGATIVE DEMAND (“CID”) SCHEDULE
FTC File No. 232-3044**

Meet and Confer: You must contact **FTC counsel**, [REDACTED] as soon as possible to schedule a telephonic meeting to be held within fourteen (14) days after You receive this CID. At the meeting, You must discuss with FTC counsel any questions You have regarding this CID or any possible CID modifications that could reduce Your cost, burden, or response time yet still provide the FTC with the information it needs to pursue its investigation. The meeting also will address how to assert any claims of protected status (e.g., privilege, work-product, etc.) and the production of electronically stored information. You must make available at the meeting personnel knowledgeable about Your information or records management systems, Your systems for electronically stored information, custodians likely to have information responsive to this CID, and any other issues relevant to compliance with this CID.

Document Retention: You must retain all documentary materials used in preparing responses to this CID. The FTC may require the submission of additional Documents later during this investigation. **Accordingly, You must suspend any routine procedures for Document destruction and take other measures to prevent the destruction of Documents in Your possession, custody, or control** that are in any way relevant to this investigation, even if those Documents are being retained by a third-party or You believe those Documents are protected from discovery. *See* 15 U.S.C. § 50; *see also* 18 U.S.C. §§ 1505, 1519.

Sharing of Information: The FTC will use information You provide in response to the CID for the purpose of investigating violations of the laws the FTC enforces. We will not disclose such information under the Freedom of Information Act, 5 U.S.C. § 552. We also will not disclose such information, except as allowed under the FTC Act (15 U.S.C. § 57b-2), the Commission’s Rules of Practice (16 C.F.R. §§ 4.10 & 4.11), or if required by a legal obligation. Under the FTC Act, we may provide Your information in response to a request from Congress or a proper request from another law enforcement agency. However, we will not publicly disclose such information without giving You prior notice.

Manner of Production: Contact [REDACTED] by email or telephone at least five days before the return date for instructions on how to produce information responsive to this CID.

Certification of Compliance: You or any person with knowledge of the facts and circumstances relating to the responses to this CID must certify that such responses are complete by signing the “Certification of Compliance” attached to this CID.

Certification of Records of Regularly Conducted Activity: Attached is a Certification of Records of Regularly Conducted Activity. Please execute and return this Certification with Your response. Completing this certification may reduce the need to subpoena You to testify at future proceedings to establish the admissibility of Documents produced in response to this CID.

Definitions and Instructions: Please review carefully the Definitions and Instructions that appear after the Specifications and provide important information regarding compliance with this CID.

I. SUBJECT OF INVESTIGATION

Whether “the “Company,” as defined herein, in connection with offering or making available products and services incorporating, using, or relying on Large Language Models has (1) engaged in unfair or deceptive privacy or data security practices or (2) engaged in unfair or deceptive practices relating to risks of harm to consumers, including reputational harm, in violation of Section 5 of the FTC Act, 15 U.S.C. § 45, and whether Commission action to obtain monetary relief would be in the public interest. See also attached resolution.

II. SPECIFICATIONS

Applicable Time Period: Unless otherwise directed, the applicable time period for the requests set forth below is from **June 1, 2020 until the date of full and complete compliance with this CID.**

A. Interrogatories

Company Background

1. State Your full legal name, principal address, other mailing addresses, principal telephone number, and date and state of incorporation; all states in which You do or have done business; and all other names under which You have done business, along with the corresponding time periods.
2. Identify Your wholly owned subsidiaries, parent companies, unincorporated divisions, joint ventures, partnerships, operations under assumed names, and predecessor companies, and describe the relationship of each to You.
3. For each year in the Applicable Time Period, state the Company’s number of employees by business unit or division and physical office location, and describe, including through graphic representation (e.g., organizational charts), how authority and responsibility were distributed among employees within the Company.
4. List all officers, directors, principals, and owners of the Company, and all shareholders with 5% or more ownership of the Company, stating each shareholder’s percentage of ownership from 2017 to present on an annual basis.
5. Describe in Detail the Company’s financial earnings, including, but not limited to, total income, gross sales and gross revenues; income/loss; and investment income, for each year in the Applicable Time Period.

6. Describe in Detail, from 2017 to the present, the Company's corporate governance model, the Company's profit cap, and any revenue sharing and cost sharing schemes or agreements.
7. Identify, from 2017 to the present, all licensing, marketing, joint venture, and research and development agreements between the Company and any other Person for any Large Language Model and/or Large Language Model Product.
8. List, by title and URL, each website or mobile application owned or operated by the Company.
9. List and Describe in Detail each Large Language Model Product (with each version of each Product listed separately) that You have provided, offered, or made available during the Applicable Time Period, including but not limited to:
 - a. The name(s) under which You have marketed, advertised, offered, or provided the Product;
 - b. The time period during which each such Product has been offered;
 - c. The language(s) in which the Product has been available and the languages in which You have promoted, marketed, or advertised the Product;
 - d. The types of Persons (e.g., individual consumers and business customers) to which each Product has been offered or made available;
 - e. A description of the functionality of the Product, including whether there are multiple tiers or versions of the Product (e.g., paid and unpaid) and, if so, the respective features or functionality of each tier;
 - f. All "plugins" available to users of the Product;¹
 - g. The Large Language Models that have been used to operate the core features of the Product, including all names by which You refer to such models or algorithms internally or publicly;
 - h. All means by which consumers have been able to access the Product, including via API or similar technology, including each website or mobile application through which You have made the Product available;
 - i. The numbers of unique individual consumer users and unique other users of the Product for each month of the applicable time period, in total and in each language in which the Product has been offered, promoted, or marketed; and

¹ OpenAI, *ChatGPT plugins*, <https://openai.com/blog/chatgpt-plugins>.

- j. The amount the Company charges for the Product and how such amount is calculated.
- 10. Identify all third parties that You have allowed use of or access to Your Large Language Models via API, and specify whether the access is paid or unpaid.
- 11. Describe in Detail how You have marketed Your Products to business customers, including by:
 - a. Identifying the top ten Persons who have purchased and/or licensed Your Large Language Models, as measured by revenue or another similar metric and, for each, Describing in Detail Your understanding of the customer's primary purpose for using Your Large Language Models; and
 - b. Describing in Detail the Company's process for identifying new business customers, including any existing or future Plans to market Your Large Language Models to specific Persons, industries, or sectors.
- 12. Identify all officers, managers, employees, or other third parties, who have prepared or supervised the preparation of responses to this CID, including their job titles, roles, and entity affiliations.

Disclosures and Representations

- 13. Describe in Detail how You have disseminated or caused to be disseminated each Advertisement, disclosure, or representation produced in response to Document Request 4 or 5, including but not limited to:
 - a. The method of dissemination;
 - b. The date(s) of dissemination;
 - c. Any steps You have taken to ensure that the Advertisement, disclosure, or representation has been clear and conspicuous to consumers;
- 14. Describe in Detail the extent to which You have conducted, caused to be conducted, relied on, or are aware of any research, copy testing, surveys, or other efforts to assess consumers' understanding of:
 - a. The methods used by Your Products as well as the accuracy or reliability of outputs generated by Your Products;
 - b. How You or third parties retain or use information collected in connection with consumers' use of Your Products; and
 - c. Your disclosures or other representations about Your Products, including those produced in response to Document Requests 4 and 5.

Model Development/Training

15. Describe in Detail the data You have used, during or prior to the Applicable Time Period, to train or otherwise develop each Large Language Model described in response to Interrogatory 9, including, for each such model:
 - a. How You obtained the data, e.g., by scraping the data, purchasing it from third parties, or by other means;
 - b. All sources of the data, including any third parties that provided data sets;
 - c. To the extent the data was derived from publicly available websites, a list of all such websites and, for each, the percentage of the data corpus that is derived from that website;
 - d. The categories of content included in the data and the extent to which each category is represented in the data corpus (i.e., as a percentage of data used to train the model);
 - e. All policies and procedures Related to identifying, assessing, vetting, and selecting sources of data for the model; and
 - f. The extent to which various languages are represented in the training data (e.g., the percentage of the training data that consists of English-language materials, Spanish-language materials, Chinese-language materials, etc.) and the extent to which the sources of data vary across languages.
16. Describe in Detail the extent to which You have, during or prior to the Applicable Time Period, assessed or reviewed the content of data used to train or otherwise develop each Large Language Model that is used in connection with providing each of Your Large Language Model Products listed in response to Interrogatory 9, including the extent to which such assessment or review has entailed the use of classifiers or other automated means and the extent to which such assessment or review has entailed manual review of the data.
17. Describe in Detail the individuals, type(s) of individuals, or departments who are responsible for training each Large Language Model Product, including:
 - a. The names of the individual(s), type(s) of individual by title and area of expertise, or department(s);
 - b. A description of their role and job functions at OpenAI; and
 - c. A description of their role(s) in training the Large Language Model, including a description of any data they created.
18. Describe in Detail the process of retraining a Large Language Model in order to create a substantially new version of the model. Include in Your response:

- a. The circumstances under which the Company opts to retrain a Large Language Model;
 - b. The individual(s), type(s) of individual, or department responsible for retraining;
 - c. The steps taken to retrain the model;
 - d. How the Company determines whether the model has been adequately retrained; and
 - e. A description of any retraining the Company actually undertook in order to correct or remediate any Large Language Model's propensity to "hallucinate"² or to reveal any Personal Information.
19. Describe in Detail the Company's process for refining a Large Language Model in order to modify an existing version of the Model. Include in Your response:
- a. The circumstances under which the Company opts to refine a Large Language Model;
 - b. The individual(s), type(s) of individual by title and area of expertise, or department responsible for and/or involved in refining;
 - c. The steps taken to refine the Model;
 - d. How the Company determines whether the Model has been adequately refined; and
 - e. A description of any refining the Company actually undertook in order to correct or remediate any Large Language Model's propensity to "hallucinate" or to reveal any Personal Information.
20. Describe in Detail the process of "reinforcement learning through human feedback."³ Include in Your response:
- a. Examples of types of experts who write questions and/or answers for use in the reinforcement learning process;
 - b. Examples of questions and answers;

² Open AI, *GPT-4*, <https://openai.com/research/gpt-4>.

³ OpenAI, *Learning to Summarize With Human Feedback*, <https://openai.com/research/learning-to-summarize-with-human-feedback>.

- c. An explanation of the rating system or method of evaluation for the Large Language Model's responses; and
 - d. A description of any reinforcement learning through human feedback the Company actually undertook in order to correct or remediate any Large Language Model's propensity to "hallucinate" or to reveal any Personal Information.
21. Describe in Detail the policies and procedures that the Company follows in order to assess risk and safety before the Company releases a new Large Language Model Product, including a new version of a Product, to the public, whether paid or for free. Include in Your response:
- a. A description of the risks taken into consideration;
 - b. A description of the individual(s), type(s) of individual by title and area of expertise, or department(s) responsible for determining whether the Large Language Model Product is ready to be released;
 - c. The role, if any, of the "deployment safety board"⁴ in the pre-release evaluation process;
 - d. A list of all instances in which a new Large Language Model Product was released to the public after this review; and
 - e. A list of all instances in which a new Large Language Model Product was *not* released to the public as a result of this review.
22. Describe in Detail the steps that the Company takes, if any, to prevent Personal Information or information that may become Personal Information when combined with other information in the training data from being included in the training data for any Large Language Model(s). Include in Your response a description of any mechanisms, processes, and/or procedures for removing, filtering, anonymizing, or otherwise obscuring such data.
23. Describe in Detail Your policies and procedures Relating to training, instructing, and overseeing human reviewers in connection with training, testing, fine-tuning, refining, monitoring, or otherwise developing Your Large Language Models or Products.

Assessing and Addressing Risks

⁴ See Nilay Patel, *Microsoft CTO Kevin Scott thinks Sydney might make a comeback*, The Verge (May 23, 2023), available at <https://www.theverge.com/23733388/microsoft-kevin-scott-open-ai-chat-gpt-bing-github-word-excel-outlook-copilots-sydney>.

24. Describe in Detail Your policies and procedures Relating to Your Large Language Model Products' actual or potential generation of statements about individuals, including but not limited to:
 - a. The extent to which the generation of statements about individuals is an acceptable or permissible outcome of users' interactions with Your Large Language Model Products;
 - b. Any conditions under which the generation of statements about individuals is viewed as an unacceptable, impermissible, or undesirable outcome of users' interactions with Your Large Language Model Products;
 - c. All policies and procedures for assessing whether the generation of statements about individuals in particular situations and scenarios is acceptable, permissible, or desirable;
 - d. All policies and procedures Related to responding to reports, complaints, or other communications reporting or alleging that Your Large Language Model Products have made statements about individuals, including false, misleading, disparaging, or harmful statements; and
 - e. The officers, managers, and business units, including, to the extent relevant, third-party contractors, with responsibilities for establishing, implementing, carrying out, and ensuring compliance with these policies and procedures, and the respective responsibilities and roles of each.
25. Describe in Detail all steps You have taken to identify or assess, test, or measure the extent to which Your Large Language Model Products have the capacity to generate statements about individuals containing real, accurate Personal Information. Your response should also include the extent to which You have conducted such testing in languages other than English.
26. Describe in Detail all steps You have taken to identify or assess, test, or measure the extent to which Your Large Language Model Products have the capacity to generate statements about real individuals that are false, misleading, or disparaging. Your response should also include the extent to which You have conducted such testing in languages other than English.
27. Describe in Detail the extent to which You have taken steps to address or mitigate risks that Your Large Language Model Products could generate statements about individuals containing real, accurate Personal Information, including but not limited to:
 - a. The extent to which You have taken steps to filter, edit, modify, supplement, or omit content from datasets used to train or otherwise develop Large Language Models, including any mechanisms, processes, and/or procedures You undertook for removing, filtering, anonymizing, or

otherwise obscuring to prevent Personal Information, or information that may become Personal Information when combined with other information in the training data, from being included in the training data for any Large Language Model(s);

- b. The extent to which You have taken steps, after the initial training or development of Your Large Language Models used to provide Your Large Language Model Products, to reinforce, reward, discourage, or otherwise increase or decrease the likelihood that Your Large Language Models will generate certain types of responses; and
 - c. All steps You have taken to filter or block Your Large Language Model Products from receiving, processing, or responding to certain types of user inputs, or to filter or block certain types of outputs, once generated by Your Large Language Models, from being delivered or displayed to users of Your Large Language Model Products.
28. Describe in Detail the extent to which You have taken steps to address or mitigate risks that Your Large Language Model Products could generate statements about real individuals that are false, misleading, or disparaging, including but not limited to:
- a. The extent to which You have taken steps to filter, edit, modify, supplement, or omit content from datasets used to train or otherwise develop Large Language Models, including any mechanisms, processes, and/or procedures You undertook for removing, filtering, anonymizing, or otherwise obscuring to prevent Personal Information, or information that may become Personal Information when combined with other information in the training data, from being included in the training data for any Large Language Model(s);
 - b. The extent to which You have taken steps, after the initial training or development of Your Large Language Models used to provide Your Large Language Model Products, to reinforce, reward, discourage, or otherwise increase or decrease the likelihood that Your Large Language Models will generate certain types of responses; and
 - c. All steps You have taken to filter or block Your Large Language Model Products from receiving, processing, or responding to certain types of user inputs, or to filter or block certain types of outputs, once generated by Your Large Language Models, from being delivered or displayed to users of Your Large Language Model Products.
29. Describe in Detail the extent to which Your Large Language Model Products have the capacity to generate statements about individuals containing real, accurate Personal Information. Your response should also include:

- a. The results of any assessment, testing, or measurement described in response to Interrogatory 25;
 - b. Any variation in those results that is correlated with the language used for testing;
 - c. Any variation in those results following or as a result of the measures described in response to Interrogatory 27; and
 - d. The evidentiary basis for Your response.
30. Describe in Detail the extent to which Your Large Language Model Products have the capacity to generate statements about real individuals that are false, misleading, or disparaging. Your response should also include:
- a. The results of any assessment, testing, or measurement described in response to Interrogatory 26;
 - b. Any variation in those results that is correlated with the language used for testing;
 - c. Any variation in those results following or as a result of the measures described in response to Interrogatory 28; and
 - d. The evidentiary basis for Your response.
31. Describe in Detail Your process for determining whether any given specific individual constitutes a “public figure” and the extent to which such determinations have any relevance to the Company’s policies, procedures, and practices Related to the development of Your Large Language models.
32. Describe in Detail the extent to which You have monitored, detected, investigated, or responded to instances in which Your Large Language Model Products have generated statements about individuals containing real, accurate Personal Information. Your response should include:
- a. The extent to which You have used classifiers or other automated means to monitor or detect the generation of statements about individuals;
 - b. The extent to which You have used manual review by human employees or contractors to review, monitor, or detect the generation of statements about individuals; and
 - c. Your policies and procedures for responding to reports, complaints, or other communications regarding specific instances in which Your Large Language Model Products have actually or allegedly generated statements about individuals.

33. Describe in Detail the extent to which You have monitored, detected, investigated, or responded to instances in which Your Large Language Model Products have generated false, misleading, or disparaging statements about individuals, including:
 - a. The extent to which You have used classifiers or other automated means to monitor or detect the generation of statements about individuals;
 - b. The extent to which You have used manual review by human employees or contractors to review, monitor, or detect the generation of statements about individuals; and
 - c. Your policies and procedures for responding to reports, complaints, or other communications regarding specific instances in which Your Large Language Model Products have actually or allegedly generated statements about individuals.
34. Describe in Detail Your policies and procedures Relating to reports, complaints, or other communications regarding specific instances in which Your Large Language Model Products have actually or allegedly generated statements about individuals containing real, accurate Personal Information (“Personal Information Complaints”), including:
 - a. For each month of the Applicable Time Period, the number of Personal Information Complaints You have received;
 - b. Captions of all lawsuits potentially Related to Personal Information Complaints;
 - c. All means by which You have received Personal Information Complaints;
 - d. Any steps You have taken to actively solicit or encourage Personal Information Complaints;
 - e. The extent to which You have conducted investigation or testing to determine whether Personal Information Complaints were accurate (e.g., whether Your Large Language Model Products had in fact generated statements about individuals), and the results of such investigations or testing; and
 - f. The extent to which You have modified Your Large Language Models or Large Language Model Products or made other technical changes in response to Personal Information Complaints.
35. Describe in Detail Your policies and procedures Relating to reports, complaints, or other communications regarding specific instances in which Your Large Language Model Products have actually or allegedly generated false, misleading,

or disparaging statements about individuals (“False Statement Complaints”), including:

- a. For each month of the Applicable Time Period, the number of False Statement Complaints You have received;
 - b. Captions of all lawsuits potentially Related to False Statement Complaints;
 - c. All means by which You have received False Statement Complaints;
 - d. Any steps You have taken to actively solicit or encourage False Statement Complaints;
 - e. The extent to which You have conducted investigation or testing to determine whether False Statement Complaints were accurate (e.g., whether Your Large Language Model Products had in fact generated statements about individuals), and the results of such investigations or testing; and
 - f. The extent to which You have modified Your Large Language Models or Large Language Model Products or made other technical changes in response to False Statement Complaints.
36. Describe in Detail the extent to which You have received complaints regarding specific instances in which Your Large Language Model Products have actually or allegedly caused any of the harms discussed as “safety challenges” in the OpenAI GPT-4 System Card,⁵ including the means by which You have received the complaints and the number of complaints You have received Related to each “safety challenge” for each Product during each month of the Applicable Time Period.

Privacy and Prompt Injection Risks and Mitigations

37. Describe in Detail the incident You publicly disclosed on or about March 24, 2020, involving “a bug in an open-source library which allowed some users to see titles from another active user’s chat history” and “payment-related information”⁶ (hereafter, “the Security Incident”), including:
- a. All dates and times during which You have reason to believe information was exposed as a result of the Security Incident;

⁵ OpenAI, *GPT-4 System Card* at 4, available at <https://cdn.openai.com/papers/gpt-4-system-card.pdf>.

⁶ OpenAI, *March 20 ChatGPT outage: Here’s what happened*, available at <https://openai.com/blog/march-20-chatgpt-outage>.

- b. All categories of information exposed;
 - c. The number of users affected, by category of information exposed;
 - d. The number of users to whom information was exposed;
 - e. The number of complaints received about the Security Incident and actions You have taken in response to those complaints;
 - f. The date on which and means by which You learned of the Security Incident;
 - g. All actions, by date, that You have taken in response to the Security Incident, including the extent to which You have provided any notifications to any users in connection with the Security Incident;
 - h. Any changes to Your policies, procedures, or practices that resulted in whole or in part from the Security Incident; and
 - i. Any testing or assessment You have conducted or relied upon in determining the extent to which there is ongoing risk to consumers in connection with the Security Incident or its underlying causes and the results of such testing or assessment.
38. Describe in Detail Your awareness of any other incident, outage, bug, or other data security vulnerability that causes or risks causing any of Your Large Language Model Products to reveal to any user any of the Personal Information, including chat history or any information about chat histories, of any other user or individual. Your response should include all subcategories of information requested in Interrogatory 37, above.
39. List all instances of known actual or attempted “prompt injection”⁷ attacks on You or Your Products, and Describe in Detail how You corrected and/or mitigated against these attacks. Include in Your response:
- a. The Product that was subject to the attack or attempted attack;
 - b. The number of users affected;
 - c. A description of how those users were affected;
 - d. How the attack was detected; and

⁷ “Prompt injection” refers to any unauthorized attempt to bypass filters or manipulate a Large Language Model or Product using prompts that cause the Model or Product to ignore previous instructions or to perform actions unintended by its developers.

- e. The source of the attack.
40. Describe in Detail all policies, procedures, mechanisms, and/or other strategies or methods for monitoring for, detecting, preventing, and/or mitigating against “prompt injection” attacks on You or Your Products.

API Integrations and Plugins

41. Describe in Detail Your policies and procedures for assessing or mitigating actual or potential risks of access to, exposure of, or exfiltration of users’ Personal Information in connection with API integrations and plugins. Include in Your response:
- a. The business units, offices, officers, or managers involved in these decisions;
 - b. The types of information and Documents that You require third parties to provide prior to their use of a plugin or API integration;
 - c. Any policies and procedures Related to testing, verifying, auditing, or otherwise monitoring plugins or API integrations. Include in Your response any general policies and procedures, and also those specific to particular plugins or API integrations;
 - d. Any technical or organizational measures You require third parties to take to assess or mitigate risks in connection with their use of Your API or plugins, including:
 - i. Restrictions on the types of third parties who may use Your API or develop plugins or the types of use cases for which third parties may use Your API;
 - ii. Restrictions on the third parties’ use of user data, including Personal Information and chat histories;
 - iii. Restrictions on types of users who may engage with the API integration or develop plugins;
 - iv. Restrictions on the individuals, categories of individual, or departments who may access user data;
 - v. Any other specific technical or organizational controls that You require third parties to implement; and
 - vi. The extent to which and means by which You monitor, enforce, or otherwise take steps to ensure that third parties comply with Your requirements, policies, and procedures Related to the use of Your API.

42. Describe in Detail any policies or procedures Related to API users' ability to "fine tune"⁸ any Model or Product using additional data or instructions. Include in Your response any measures You have taken to ensure that "fine tuning" does not increase risks of unauthorized access to, exposure of, or exfiltration of users' Personal Information or that Your Large Language Models may generate statements about individuals, including false or derogatory statements.

Monitoring, Collection, Use, and Retention of Personal Information

43. For each Product, state which individuals, type(s) of individual by title and area of expertise, or departments may access or monitor user interactions with the Product, whether that access or monitoring is real-time or later in time.
44. For each individual, category of individual, or department You included in response to Interrogatory 43, state the purpose(s) for which user interactions are accessed or monitored.
45. For each Product, list all the types of Personal Information collected, used, analyzed, stored, or transferred by the Company during the Applicable Time Period. For each type of Personal Information listed, Describe in Detail:
- a. The source of the data;
 - b. Where each type of data is stored, including the extent to which You have created copies of the data for product development or other purposes;
 - c. How long the data is stored by default;
 - d. How long the data is stored if a user has opted out of data retention;
 - e. How long the data is stored after a user has requested its deletion or has deleted their account;
 - f. The purpose(s) for which the Company collects, uses, analyzes, stores, or transfers the data;
 - g. The individuals or categories of individuals who can access the data; and
 - h. The purpose(s) for which those individuals or categories of individual can access the data.
46. For each Product, Describe in Detail each and every mechanism or procedure available to users or other individuals (including users of third-party licensees' products or services) to opt not to have the Company collect, retain, use, analyze,

⁸ "Fine tune" or "fine tuning" has the same meaning as in this document disseminated on Your website at <https://platform.openai.com/docs/guides/fine-tuning>.

transfer, and/or otherwise access individuals' Personal Information (each an "opt-out"). For each mechanism or procedure, describe:

- a. How the request is effectuated;
 - b. The timeframe in which the request is effectuated;
 - c. The data to which the opt-out applies;
 - d. Which data the Company collects notwithstanding the opt-out;
 - e. The circumstances under which the Company collects, retains, uses, analyzes, transfers, and/or otherwise accesses that data, notwithstanding the opt-out;
 - f. Where that data is stored;
 - g. How long that data is stored;
 - h. The purpose for which the Company collects, retains, uses, analyzes, transfers, and/or otherwise accesses that data; and
 - i. Any process by which users can verify the efficacy of their opt-out.
47. For each Product, Describe in Detail each and every data store where Personal Information is stored or used, whether temporarily or permanently. For each data store described, include:
- a. The name of the data store, including the name used colloquially within the Company;
 - b. The storage capacity of the data store;
 - c. The primary function of the data store;
 - d. The types and/or sources of data stored in the data store;
 - e. The individuals or categories of individuals who have access to the data store;
 - f. The circumstances under which that data store's data is copied, removed, transferred, modified, or otherwise accessed;
 - g. The security measures, whether physical or technological, protecting or intended to protect the security of the data in the data store; and
 - h. Any and all policies and procedures Relating to mapping or tracking how data is copied and stored, and maintaining the security of the data store and/or the privacy of the data contained in it.

48. For each Product, Describe in Detail each and every mechanism or procedure available to users to request that the Company delete any of their Personal Information, including chat histories. For each mechanism or procedure, describe:
- a. How the request is effectuated;
 - b. The timeframe in which the request is effectuated;
 - c. The data to which the deletion request applies;
 - d. Which data the Company retains notwithstanding the deletion request;
 - e. The circumstances under which the Company retains that data, notwithstanding the deletion request;
 - f. Where that data is stored;
 - g. How long that data is stored;
 - h. The purpose for which the Company retains, analyzes, transfers, and/or otherwise accesses that data; and
 - i. Any process by which users can verify the efficacy of their deletion request.
49. For each Product, Describe in Detail any efforts the Company undertakes to aggregate, anonymize, pseudonymize, or otherwise de-identify Personal Information collected from users. Include in Your response:
- a. How the data was aggregated, anonymized, pseudonymized, or de-identified;
 - b. Whether such aggregation/anonymization/pseudonymization/de-identification was accomplished by a third party or by You;
 - c. The fields or categories of Personal Information that are contained in the data prior to aggregation, anonymization, pseudonymization, or de-identification;
 - d. The fields or categories of Personal Information that are contained in the data after aggregation, anonymization, pseudonymization, or de-identification; and
 - e. Whether aggregated, anonymized, pseudonymized, or de-identified data is stored separately from other data, and if so, where.

B. Requests for Documents

1. Produce organizational charts sufficient to show how responsibilities have been allocated among officers, managers, and business units of the Company throughout the applicable time period.
2. Produce all contractual agreements of the following types offered, requested, or signed, from 2017 to the present, between the Company and any Person Related to Your Large Language Model and/or Large Language Model Products:
 - a. Any partnership agreements, joint venture agreements, revenue sharing or cost sharing agreements, and any agreements containing an exclusivity term; and
 - b. All licensing, marketing, and research and development agreements;
3. For each contractual agreement produced in response to Document Request 2, produce:
 - a. All Documents exchanged from 2017 to the present between You and any Person Relating to the contractual agreement; and
 - b. All Documents shared from 2017 to the present with Your Board of Directors or other senior-level executives discussing the negotiations of the agreement.
4. Produce all Advertisements or public statements that You have disseminated or caused to be disseminated Related to the capabilities, accuracy, or validity of the output of any Product or actual or potential use cases for any Product.
5. Produce all representations or disclosures that You have disseminated or caused to be disseminated regarding the limitations of or risks associated with any of Your Products, including limitations or risks Relating to the potential to produce inaccurate statements or risks to the privacy or security of consumers' Personal Information.
6. Produce all Documents Related to any research, copy testing, surveys, or other efforts to assess consumers' understanding of Your products and services, including with respect to:
 - a. The accuracy or reliability of outputs generated by Your Products;
 - b. How You or third parties retain or use information collected in connection with consumers' use of Your Products; and
 - c. Your disclosures or other representations about Your Products.
7. Produce all contracts, including any amendments, that You have entered into with third parties, prior to or during the Applicable Time Period, in connection with obtaining data that You have used to train one or more Large Language Models

that You have offered, provided, or made available during the Applicable Time Period.

8. Produce all contracts, including any amendments, that You have entered into with third parties in connection with their operation of API integrations of Your Large Language Models or “plugins,” including all contracts with Persons Identified in response to Interrogatories 10-11.
9. Produce all Documents Relating to the testing, assessment, or measurement described in response to Interrogatories 25, 26, 35, and 37 above, including the results of all such testing.
10. Produce all internal communications about or assessments of Your Products’ potential to
 - a. Produce inaccurate statements about individuals; or
 - b. Reveal Personal Information.
11. Produce all Documents Relating to policies and procedures described in response to any Interrogatory contained herein.
12. Produce Documents sufficient to show Your policies and procedures Relating T=to the privacy and security of Personal Information.
13. Any policies, guidance, dictionaries, or other written guidance defining or otherwise explaining OpenAI’s interpretation of the following terms as used in Your public statements:
 - a. “freely and openly available” data;⁹
 - b. “reinforcement learning from human feedback;”¹⁰ and
 - c. “hallucination” or “hallucinate,” as it pertains to a Large Language Model’s propensity to generate responses that are false or unjustified.¹¹
14. Produce all instructions and training materials provided to all employees, contractors, volunteers, or other humans who label, rate, review, and/or otherwise evaluate any of Your Large Language Models’ responses to user inputs or queries.

⁹ OpenAI, *How ChatGPT and Our Language Models Are Developed*, <https://help.openai.com/en/articles/7842364-how-chatgpt-and-our-language-models-are-developed>.

¹⁰ OpenAI, *Learning to Summarize With Human Feedback*, <https://openai.com/research/learning-to-summarize-with-human-feedback>.

¹¹ Open AI, *GPT-4*, <https://openai.com/research/gpt-4>.

15. Produce Documents sufficient to show examples of the types of prompts or instructions You use to retrain your Large Language Models.
16. Produce, consistent with Instruction I-11, sample (to be prescribed by FTC counsel at a later date) complaints, reports, or communications from consumers, users, or other third parties Relating to:
 - a. Instances in which Your Large Language Model Products have generated statements about individuals containing accurate Personal Information;
 - b. Instances in which Your Large Language Model Products have generated inaccurate or disparaging statements about individuals;
 - c. The privacy, confidentiality, security, or integrity of consumers' Personal Information; and
 - d. Prompt injection attacks.
17. Produce, consistent with Instruction I-11, sample (to be prescribed by FTC counsel at a later date) consumer complaints regarding specific instances in which Your Large Language Model Products have actually or allegedly caused any of the harms discussed as "safety challenges" in the OpenAI GPT-4 System Card, including generation of harmful or misleading content.

NOTICE: This CID does not seek any information that is prohibited from disclosure under the Cable Communications Policy Act of 1984 ("Cable Act"), 47 U.S.C. §§ 551 et seq., the Satellite Television Extension and Location Act ("STELA"), 47 U.S.C. § 338(i), or the Electronic Communications Privacy Act ("ECPA"), 18 U.S.C. §§ 2701 et seq. To the extent that You are, for purposes of ECPA, a provider of Electronic Communications Service or Remote Computing Service to a customer or subscriber about whom this CID seeks information, do not divulge a record or information pertaining to such customer or subscriber or the content of such customer's or subscriber's communications, other than the content, records, and information specifically requested in this CID. If You have any questions, please contact FTC counsel before providing responsive information.

III. DEFINITIONS

The following definitions apply to this CID:

D-1. **"Advertisement"** or **"Advertising"** or **"Ad"** means any written or verbal statement, illustration, or depiction that promotes the sale of a good or service or is designed to increase consumer interest in a brand, good, or service. Advertising media includes, but is not limited to: packaging and labeling; promotional materials; print; television; radio; and Internet, social media, and other digital content.

D-2. **"API"** or **"Application Programming Interface"** means a source code-based specification intended to be used as an interface by software components to communicate with each other.